

(19)

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 1 026 582 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:

09.08.2000 Bulletin 2000/32

(51) Int. Cl.⁷: G06F 9/38

(21) Application number: 00300821.6

(22) Date of filing: 02.02.2000

(84) Designated Contracting States:

AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE

Designated Extension States:

AL LT LV MK RO SI

(30) Priority: 04.02.1999 US 243982

(71) Applicant:

SUN MICROSYSTEMS, INC.
Palo Alto, California 94303 (US)

(72) Inventor: Harris, Jeremy G.

Buckinghamshire HP8 4AF (GB)

(74) Representative:

Haines, Miles John et al

D. Young & Co.


21 New Fetter Lane

London EC4A 1DA (GB)

(54) Handling of load errors in computer processors

(57) In a RISC or CISC processor supporting the IEEE 754 Not-a-Number (NaN) standard and of the kind comprising a load/store unit, a register unit and an arithmetic logic unit, and wherein the load/store unit has an error flag for marking a datum loaded to the load/store unit following a load which has completed, but resulted in an error, the processor is provided with a bit pattern generator operatively arranged in an output path from the load/store unit to at least one of the register unit and the arithmetic logic unit so that a Not-a-Number value for the invalid datum is loaded into a destination one of the floating-point registers or the arithmetic logic unit. The arithmetic logic unit is configured to propagate the Not-a-Number value as a Quiet-Not-a-Number (QNaN) value through its operations. The QNaN value may be tested for in a datum by a system software command code provided for that purpose. In other embodiments of the invention, similar functionality is provided for integer registers using poison/valid bits in conjunction with an arithmetic logic unit designed to propagate the poison/valid bits through its operations. An advantage to be gained by this design is that it becomes possible to delay testing the results of a non-faulting load, since the QNaN-like symbolic entity will propagate with the results of operations on an invalid datum, thereby keeping track of the integrity of the data.

Fig. 1



EP 1 026 582 A2

Description

BACKGROUND OF THE INVENTION

[0001] The invention relates generally to the handling of load errors in computer processors, more especially but not exclusively to the handling of load errors resulting from speculative loads.

[0002] For good performance of a processor in a data processing system it is desirable to overlap data loads with other operations, by moving the load instructions forward to earlier positions in the instruction stream. When a load instruction is moved ahead of conditional control structures in the program flow, then the address it reads from may not yet be validated by the rest of the program code and may therefore be wrong. Loading of this kind is referred to as speculative.

[0003] A speculative load is thus defined a load operation that is issued by a processor before it is known whether the results of the load will be required in the flow of the program. Speculative loads can reduce the effects of load latency by improving instruction scheduling. Generally, speculative loads are generated by the compiler promoting loads to positions before test control instructions.

[0004] Speculative loads are often implemented as non-faulting loads. A non-faulting load is a load which always completes, even in the presence of faults. The semantics of a non-faulting load are the same as for any other load, except when faults occur. An example of a fault is an address-out-of-range error. When a fault occurs, it is ignored and the hardware and system software cooperate to make the load appear to complete normally, but in some way to return the result in a form which reflects that the loaded datum is invalid. Typically, the hardware will be configured to generate a fault indication for a failed normal load and to return a particular data value for a failed speculative load.

[0005] One known example of the handling of a failed load is the standard use of a poison bit or valid bit in the register into which the result of the speculative load is loaded. If the non-faulting load is successful then the poison bit remains unset or the valid bit is set. On the other hand, if the non-faulting load is unsuccessful then the poison bit is set or the valid bit remains unset. The software or hardware is then configured to ensure that any subsequent use of the data in the register generates a trap. With this approach, whenever there is an error in a non-faulting load, the program flow will enter into an error handling routine when an operation on the invalid data is attempted.

[0006] Another example of the handling of a failed load is to be found in the Sun SPARC processors UltraSPARC I & II. Here a non-faulting load returns zero-valued data when an exception (i.e. an error) is encountered. Software code then uses a compare instruction to check the load result before use, not using the speculatively loaded data if it is zero. If the result is

zero, then the memory address is read again later using a normal (non-speculative) load to which normal protection mechanisms apply. The normal load will be able to differentiate between correct zero-valued data and an exception condition. Only if the normal load shows an error will an exception be caused, i.e. a trap generated. With this approach, whenever a zero result is returned from a non-faulting, speculative load, the instruction stream is stalled until the normal load has completed.

[0007] It is an aim of the invention to provide a mechanism for handling non-faulting loads which can improve program flow in the cases that non-faulting loads return invalid results.

SUMMARY OF THE INVENTION

[0008] Particular and preferred aspects of the invention are set out in the accompanying independent and dependent claims. Features of the dependent claims may be combined with those of the independent claims as appropriate and in combinations other than those explicitly set out in the claims.

[0009] According to a first aspect of the invention there is provided a processor comprising a load/store unit, a register unit comprising a set of registers and an arithmetic logic unit, the processor being of the kind in which the load/store unit has an error flag for marking as invalid a datum loaded to the load/store unit following a load which has not reliably completed and which is thus to be treated as having failed. The processor is modified by the provision of a symbolic entity transmitter operatively arranged as an output stage of the load/store unit so that a symbolic entity is loaded into a destination one of the registers or directly into the arithmetic logic unit when the error flag is set in the load/store unit following a failed load. Moreover, the arithmetic logic unit is configured to propagate the symbolic entity, when present in an operand of an operation carried out by the arithmetic logic unit, to a result of the operation, the result with symbolic entity then being conveyed either to a destination register of the register unit or the load/store unit, depending on the processor design.

[0010] In the present document, it should be noted that the term arithmetic logic unit (ALU) is used as a generic term for both integer logic units (which in the art are usually referred to as arithmetic logic units) and floating point units (FPUs).

[0011] In the case of floating-point registers in a processor conforming to IEEE 754, the symbolic entity may be a Not-a-Number (NaN) value. The symbolic entity transmitter may then take the form of a bit pattern generator interposed between the load/store unit and the register unit, and/or between the load/store unit and the ALU. The bit pattern generator is then configured and arranged to load a bit pattern of a NaN value into the load destination register or the ALU in the case of a failed load. The NaN value may be one of the large number of defined NaN values which is not used as a

NaN value by the remaining hardware. Alternatively, a NaN value used by the processor for other purposes may be used. No or minimal special hardware is required in the ALU, since the ALU will automatically propagate a NaN value through arithmetic and logical operations. Moreover, no additional internal bandwidth will be required for the communication links between the load/store unit, register unit and ALU, since the failed load information is conveyed with the normal data bits.

[0012] Thus, according to a floating-point aspect of the invention, there is provided a bit pattern generator operatively arranged in an output path from the load/store unit so that a Not-a-Number value for the invalid datum is loaded into a destination one of the floating-point registers in the register unit, or directly into the ALU.

[0013] The ALU is preferably configured to propagate the Not-a-Number value as a Quiet-Not-a-Number (QNaN) value through operations carried out in the ALU. Moreover, the QNaN value is preferably testable for in a datum by a system software command code provided for that purpose. The command code may include a conversion, conditional on the test result, of the QNaN value to a Signaling-Not-a-Number (SNaN) value, so as to cause generation of a trap on subsequent use of the datum concerned. This may be especially useful in a processor supporting multiple threads of control where much of the processor execution will involve computing alternative "ways" only one of which will ultimately lie on the execution path of the code. Alternatively, the command code may include a conditional branch, conditional on the test result, for immediately invoking an error handling routine for dealing with the invalid datum.

[0014] In the case of integer registers, the symbolic entity transmitter may take the form of hardware interposed between the output-side of the load/store unit and the input-side of the register unit and/or ALU so that, in the case of a failed load, the error flag set in the load/store unit is conveyed to set or unset a poison or valid bit, respectively, in the destination register or operand of the ALU. Moreover, the ALU and register unit, or ALU register unit and load/store unit, are interconnected so as to transmit and receive the poison or valid bit from each other during processor operation, and the ALU is internally configured to propagate the poison or valid bit through its operations.

[0015] Thus, according to an integer embodiment of the invention, there is provided a reduced instruction set computer (RISC) processor, the register unit of which includes a set of integer registers, the integer registers having one or more poison bits responsive to loads from the load/store unit of invalid data, or, alternatively, one or more valid bits responsive to loads from the load/store unit or ALU of valid data, the poison or valid bits thus serving to indicate the integrity of data held in the respective registers, wherein the ALU is configured to propagate poison or valid bits present in operands of the operations to the results of the operations and to

return the results together with the propagated poison or valid bits to the registers or ALU. It is noted that poison or valid bits may also be used as the symbolic entities for floating point data instead of Not-a-Number values.

[0016] Similar functionality may be provided in an another integer embodiment of the invention in a complex instruction set computer (CISC) processor by configuring the ALU to receive operands including one or more poison bits responsive to loads from the load/store unit or register unit of invalid data, or, alternatively, one or more valid bits responsive to loads from the load/store unit or register unit of valid data.

[0017] When performing an operation, the ALU of a processor according to the above-described integer aspect of the invention will return a poisoned or non-valid result if one or more of the operands of the operation are poisoned or non-valid respectively. On the other hand, the ALU will, in all cases, or all but a number of special cases, return a valid or non-poisoned result if the or each operand of the operation is valid or non-poisoned respectively. The special cases where a poisoned or non-valid result will be returned even when all operands are non-poisoned or valid will be those in which the result of an operation can be predicted as being invalid merely by virtue of, first, the operation type and, second, either the value of one operand or the combination of values of two or more operands.

[0018] A software command code may be provided for testing a datum for the presence of poison or valid bits. Moreover, a branch conditional on the result of the testing may form part of the software command code execution, whereby an invalid datum can be handled by branching to an error handling routine.

[0019] According to a further aspect of the invention there is provided a method of operation of a processor comprising an instruction unit, a load/store unit, a register unit comprising a set of registers, and an ALU, the method comprising the steps of:

- a) the instruction unit issuing a load request for a datum to an external storage element;
- b) the load being carried out, but returning an invalid datum to the load/store unit;
- c) the load/store unit setting an error flag for the invalid datum;
- d) the load completing by loading a symbolic entity as at least a part of the datum into one of the register unit and the arithmetic logic unit;
- e) the arithmetic logic unit carrying out an operation having the datum as an operand such that the symbolic entity associated with the invalid datum is conveyed to a result of the operation; and
- f) outputting from the arithmetic logic unit the result with symbolic entity into one of the register unit and the load/store unit.

[0020] The invention, especially in some of its

embodiments for floating-point operations, may be better understood by analogy with the concept of Not-a-Number (NaN) as defined by IEEE 754 (1985) which is a well-known standard for floating-point arithmetic implemented in many recently designed processors.

[0021] NaN is described in IEEE 754 and in standard literature on the programming of any processor which implements NaN. A brief summary is however now given. NaN is a symbolic entity encoded in floating-point format. The IEEE floating-point single and double formats includes a sign bit, a number of exponent bits, in the form of a biased exponent, and a number of mantissa bits conveying the fraction. The sign and mantissa collectively form the significand. Reserved values of the exponents are used to encode NaN's. The reserved values may be any values apart from those two reserved for \pm infinity. If the biased exponent is all ones (in its binary representation) and the fraction is not zero then the significand conveys a NaN.

[0022] The NaN standard applies to arithmetic operations such as add, subtract, multiply, divide and square root, as well as to various other arithmetic and logical operations, such as conversions between number formats, remaindering and rounding, and optionally copying without change of format. If one or more signaling NaN (SNaN) values are input to an operation then an exception is signalled. If one or more quiet NaN (QNaN) values are input to an operation, and no SNaN's, then the operation signals no exception and delivers as its result a QNaN. With each exception there is typically an associated trap under software control. SNaN's thus signal the invalid operation exception whenever they appear as operands and will result in the setting of a status flag, taking a trap or both. On the other hand, QNaN's propagate through almost every arithmetic operation without signaling exceptions.

[0023] Now, by analogy with NaN, the invention may be thought of as the provision of a symbolic entity similar to a propagating QNaN which propagates when operations are carried out on an invalid datum resulting from a failed speculative load.

[0024] In contrast, with fore-knowledge of the present invention, the prior art use of a poison bit or valid bit in conjunction with trapping, as described above, may be thought of as analogous to the provision of a non-propagating SNaN. A QNaN-type functionality cannot be provided with the conventional design, since there is no means for generating a propagatable QNaN-like bit pattern in the load/store unit when the flag for a failed speculative load is set. Moreover, for non-floating-point operations, a conventional ALU has no means for propagating a QNaN-like entity through its operations, even if one could be generated.

[0025] An advantage achievable with some embodiments of the invention is that it becomes possible to delay testing the results of a non-faulting load, since the QNaN-like symbolic entity will propagate with the results of operations on an invalid datum, thereby keep-

ing track of the integrity of the data. This has the benefit that program flow does not have to be slowed or otherwise disrupted by testing non-faulting load results as the loads occur, but can be deferred until some other time, for example when processor time is freely available, under the control of the program. The testing can be performed when convenient and, if the test reveals a QNaN-like entity, i.e. data corrupted by an earlier failed speculative load, then this can be dealt with, for example immediately by branching to a servicing routine, or by convening the QNaN-like entity to a SNaN-like entity so as to cause trapping on subsequent use.

[0026] Although the invention was conceived with the handling of failed speculative loads specifically in mind, it will be appreciated that the processor designs herein described are equally well suited to handling failed loads of any type. Delay of testing the results of loads of any kind may be advantageous for the same reason as described above in relation to non-faulting loads. One example of the utility of the propagatable, error indicating, symbolic entity for normal loads would be to provide hardware protection against programming errors of the kind which may result in illegal loads from external storage, e.g. a load from an address that does not exist. The invention may thus be embodied in processors that do not support non-faulting loads, and also in processors that do support non-faulting loads for handling both failed normal loads and failed speculative loads.

BRIEF DESCRIPTION OF THE DRAWINGS

[0027] For a better understanding of the invention and to show how the same may be carried into effect reference is now made by way of example to the accompanying drawings in which:

Figure 1 shows a floating-point processor according to a first embodiment of the invention;

Figure 2 shows the format of a floating-point number supported by the processor of Figure 1;

Figure 3 is a table of the exponent, significand and fraction parts of the floating-point format of Figure 2;

Figure 4 shows an integer processor according to a second embodiment of the invention;

Figure 5 shows the flow of an operation supported by the arithmetic logic unit of the integer processor of Figure 4;

Figure 6 shows a combined floating-point and integer processor according to a third embodiment of the invention;

Figure 7 shows the flow of conversion operations supported by the arithmetic logic unit of the processor of Figure 6.

Figure 8 shows a floating-point processor according to a sixth embodiment of the invention; and

Figure 9 shows an integer processor according to a

seventh embodiment of the invention.

DETAILED DESCRIPTION

[0028] Figure 1 shows a processor according to a first embodiment of the invention. The processor comprises the conventional components of an instruction unit 1, a load/store unit 2, a register unit including a set of floating-point registers 4, and an arithmetic logic unit (ALU) 5 in the form of a floating point unit (FPU). These components are interconnected internally and externally in a generally conventional fashion for a reduced instruction set computer (RISC) processor. Namely, there is an internal line of control from the instruction unit 1 to the load/store unit 2, from the instruction unit 1 to the ALU 5, and from the load/store unit 2 to the register unit 4, and also bi-directional internal lines of control between the register unit 4 and the ALU 5. Furthermore, external to the processor unit there is a memory cache 3 which interfaces with the processor unit through the load/store unit 2 through an external input/output of the load/store unit 2 provided for that purpose. This arrangement is also conventional. Instead of a cache, the processor may interact directly with any other external storage, such as a main memory unit.

[0029] The floating-point processor of Figure 1 is of the kind capable of supporting speculative loads from an external data storage, in this case the cache 3. For this purpose, the load/store unit 2 has an error flag for marking a datum loaded to the load/store unit 2 following a speculative load which has loaded a datum that has not yet been validated by the rest of the program code and may therefore be wrong.

[0030] The generally conventional RISC architecture of the processor unit of Figure 1 is departed from by the inclusion of a bit pattern generator 6 operatively arranged on the internal output side of the load/store unit 2, more particularly in the data path from the load/store unit 2 to the register unit 4. The reverse path from the register unit 4 to the load/store unit 2 is not affected by the bit pattern generator 6. The bit pattern generator 6 is configured to generate as an output a Not-a-Number (NaN) value. In the course of a load completion, a datum is loaded from the load/store unit 2 into a destination one of the floating-point registers contained in the register unit 4. The bit pattern generator 6 is responsive to the error flag in the load/store unit 2 for the datum being loaded. If the error flag is not set, then the datum is passed unchanged from the load/store unit 2 to the register unit 4, i.e. the bit pattern generator 6 has no operational effect. On the other hand, if the error flag in the load/store unit 2 is set, indicating a load which has not completed reliably, i.e. failed, then the bit pattern generator 6 acts to load the reserved NaN value into the destination register, instead of the datum loaded into the load/store unit 2 from the external storage 3. The error flag will be set following a failed load of any kind and a failed speculative load is but one example of

a load error.

[0031] The ALU 5 is configured to handle the NaN value generated by the bit pattern generator 6 as a Quiet-not-a-Number (QNaN) value. As described further above, a QNaN value is one which propagates through operations in the ALU 5 without signaling exceptions. The NaN value generated by the bit pattern generator 6 may be a general QNaN value used for normal purposes in the processor, or may be a separate NaN value which the processor is also designed to propagate in the manner of a QNaN and which is specific to data resulting from failed loads. (Here it is noted that current processors typically has a very large number of unused NaN values. For example, in the Intel pentium processor for single precision format floating-point numbers there are 2^8-3 available NaN values only a few of which are used).

[0032] In the case that a general QNaN value is generated by the bit pattern generator, the qualities of the conventional floating-point NaN value would be merged with the non-valid-datum NaN value. The special NaN value usage according to the invention would then be integrated with the semantics required by conventional NaN value handling of a processor that conforms with IEEE 754.

[0033] On the other hand, in the case that a separate QNaN value is designated specifically for the purpose of tagging non-validated data resulting from failed loads, there is then the possibility of providing a more sophisticated protocol for responding to failed loads, either programmatically or in hardware. Furthermore, different QNaN values may be used for different types of load error. For example, data from failed speculative loads may be ascribed one QNaN value and data from failed normal loads another QNaN value.

[0034] Figure 2 shows the basic format of a floating-point number, for example either in conventional single or double precision, as used in IEEE 754. The floating-point numbers are composed of three fields, namely a one bit sign field "s", a multi-bit biased exponent "e" and a multi-bit fraction "f". In conventional single precision format, the exponent has 8 bits and the fraction 23 bits, thereby to provide a significand of 24 bits (including the sign bit). Conventionally, and as shown in Figure 3, NaN values are defined as values of the fraction part of the significand which may have any value apart from 10...00 when all the bits of the biased exponents are ones, i.e. 11...11, and when the sign bit has a negative value.

[0035] In operation, the bit pattern generator 6 thus ensures that speculative loads which load invalid, or non-validated, data will result in QNaN values being loaded into the destination registers of the register unit 4. The conventional NaN value support provided in the ALU 5 will then ensure that any results of operations which use any of the invalid data are themselves given the QNaN value. In this way, program flow need not be interrupted until a convenient time, since data integrity is tracked through the QNaN propagation process. The

particular action needed to deal with a failed load may be delayed for as many instructions as needed, while still permitting manipulation and use of the results of the load, should it complete normally in the manner of a non-faulting load.

[0036] In one implementation, the system software code is provided with a command for testing for the validity of any datum by means of testing the datum for the QNaN value generated by the bit pattern generator 6. Moreover, the processor is configured, or its software designed, so that the QNaN test command is required to be issued before selected operations, such as an attempt to use a non-valid datum as an address for a load or store, an attempt to store a non-valid datum, or an attempt to use a non-valid datum to alter internal processor functional modes. A positive test result may then cause a trap to be generated or a branch to a servicing routine. A still further option would be to provide a trap on use of a non-valid datum, the trap being switchable by the program.

[0037] A trapping function may be implemented by providing a command for converting the value of a datum to a SNaN value. The SNaN value may be either a general SNaN value used in normal processor operation, or one or more special SNaN values specific to data resulting from failed loads. (Here it is noted that any processor supporting IEEE 754 will have at least one SNaN value and associated support). The conversion command would typically be used following a test command returning a positive result, thereby converting a QNaN value into a SNaN value. The SNaN value will automatically trap on use, or invoke whatever other general SNaN servicing routines are present in the processor hardware or software, since, by definition, a SNaN value will signal an invalid operation exception whenever it appears as an operand of an operation.

[0038] Figure 4 shows a processor according to a second embodiment of the invention interfaced to an external storage element. The processor unit of the second embodiment includes an instruction unit 1, a load/store unit 2, a register unit 4 including a set of integer registers, and an ALU 5 for integer arithmetic operations. The processor unit is connected to an external storage element in the form of a memory cache 3, the processor unit being connected to the cache 3 through the load/store unit 2. The components of the processor unit are connected in a generally conventional manner. Namely, there is a line of control from the instruction unit 1 to the load/store unit 2, from the instruction unit 1 to the ALU 5, and from the load/store unit 2 to the register unit 4, and also bi-directional lines of control between the register unit 4 and the ALU 5.

[0039] By contrast to a conventional RISC processor unit, the communication path between the register unit 4 and the ALU 5 is provided with one bit of extra bandwidth through an additional bus line 7 which is shown in Figure 4, for the sake of convenience only, as being separate from the remaining internal bus lines

interconnecting the register unit 4 and the ALU 5. An extension of the bus line 7 is linked to the load/store unit 2.

[0040] Each of the integer registers of the register unit 4 have a poison/valid bit responsive to loads from the load/store unit 2 of failed, speculatively loaded, data. The poison/valid bits thus indicate the integrity of the datum held in the integer register concerned. The provision of poison/valid bits responsive to the error flag in the load/store unit is known, as is the loading of a poison or valid bit into a destination register of the register unit from the load/store unit. However, the processor of Figure 4 differs from a conventional design in that the register unit 4 and ALU 5 are configured to pass the poison/valid bits with their associated data during intercommunication between these components. Moreover, the ALU 5 is configured to propagate the poison/valid bits present in any operand of an operation to the result or results of that operation, and to return the results together with the propagated poison/valid bits to the integer registers in the register unit, using the additional bus line 7 for communication of the poison/valid bits. In this way, the same functionality as described above for the first embodiment in respect of floating-point registers, i.e. the quiet propagation of a symbolic entity tagging data validity, can also be provided for integer registers.

[0041] As shown in Figure 5, a given operation carried out in the ALU 5 will have a first integer operand 8, and optionally one or more further integer operands 9. Each operand includes an integer datum part "i" and a poison bit "p". (Here it will be understood that valid bits may be used instead of poison bits and also that multiple poison bits or valid bits may be used instead of a single poison bit or a single valid bit respectively). The operation on the or each operand will complete to generate an operation result 10, also having an integer datum portion "i" and a poison bit "p". The hardware of the ALU 5 is configured to ensure that the poison bit of the result 10 is set if one or more of the poison bits of the or each operand is set. As will be appreciated, the internal hardware of the ALU 5 may be implemented to provide this functionality for all of its operations or only for a defined subset of its operations.

[0042] Figure 6 shows a third embodiment of the invention which is a hybrid of the first and second embodiments of the invention in which the register unit 4 has a set of floating-point registers and a set of integer registers, both of which support propagation of a data-integrity signifying symbolic entity. A bit pattern generator 6, as described in connection with the first embodiment, is provided to support the functionality for the floating-point registers and a bus line 7 and modified internal hardware design of the ALU 5 is provided to support the same functionality for the integer registers, as in the second embodiment.

[0043] The ALU of the third embodiment supports both floating point and integer arithmetic operations.

Floating-point operations are handled in the manner described for the first embodiment and integer operations are handled in the manner described for the second embodiment. In addition to these operations, there will also be integer-to-floating-point conversion operations in the third embodiment, as shown in Figure 7 with the solid arrow. Such a conversion operation has an integer operand. The conversion process is responsive to the poison bit which, if set, will result in the reserved QNaN value generated by the bit pattern generator 6 being defined as the result of the conversion process, as shown in Figure 7. Similarly the ALU 5 is configured to support the reverse floating-point-to-integer conversion process, setting the poison bit of the result if the operand has the QNaN value, or indeed any NaN value, as shown in Figure 7 by the dashed arrow. The ALU 5 is also configured to set the poison bit in the case of a conversion overflow.

[0044] According to a fourth embodiment of the invention (not separately illustrated) there is provided a floating point processor having the same basic structure as shown for the integer processor of Figure 4. In the floating point processor of the fourth embodiment, poison/valid bits are used as the symbolic entities indicating an error status in the datum concerned, as opposed to NaN-like symbolic entities, as in the first embodiment. The load/store unit of the fourth embodiment may optionally be provided with a conversion unit. The conversion unit is operable on storing a datum to convert a datum with set poison or unset valid bits to a NaN value. Moreover, the conversion unit is operable on loading a datum to set the poison or unset the valid bits when a datum having a NaN value is loaded into the processor.

[0045] According to a fifth embodiment of the invention (not separately shown) there is provided a combined integer and floating point processor having the same basic structure as shown for the integer processor of Figure 4. In the processor of the fifth embodiment, poison/valid bits are used as the symbolic entities both for integer data and floating point data, substantially as described with reference to the fourth embodiment.

[0046] The above described embodiments relate to processor units with a reduced instruction set computer (RISC) architecture. The invention may however also be embodied in processor units with a complex instruction set computer (CISC) architecture, as described in more detail further below. In a CISC architecture, data which are operands of operations to be carried out in the ALU may be directly loaded from the load/store unit into the ALU, not via the register unit. Moreover, the ALU may directly output data, for example the results of operations, to the load/store unit not via the register unit. The register unit is connected to the ALU, as in a RISC architecture, and can thus serve to store intermediate results from the ALU, for example. CISC architectures may also allow for loading of data from the load/store unit into the register unit and storing of data from the register unit to the load/store unit, as in a RISC architec-

ture.

[0047] Figure 8 shows a processor according to a sixth embodiment of the invention. The processor comprises the conventional components of an instruction unit 1, a load/store unit 2, a register unit including a set of floating-point registers 4, and an ALU 5 in the form of a FPU. These components are arranged in a generally conventional fashion for a complex instruction set computer (CISC) processor. Namely, there is a line of control from the instruction unit 1 to the load/store unit 2 and from the instruction unit 1 to the ALU 5, and also bi-directional lines of control between the load/store unit 2 and the ALU 5 and between the register unit 4 and the ALU 5. In addition, as shown by dashed lines, bi-directional lines of control may exist between the load/store unit 2 and the register unit 4. Furthermore, external to the processor unit there is a memory cache 3, or other external storage, which interfaces with the processor unit through the load/store unit 2, more especially though an input/output arranged on the external side of the load/store unit 2. This arrangement is also conventional.

[0048] The floating-point processor of Figure 8 is of the kind capable of supporting speculative loads from an external data storage, in this case cache 3, using an error flag as described for the first embodiment.

[0049] The generally conventional CISC architecture of the processor unit of Figure 8 is departed from by the inclusion of a bit pattern generator 6 operatively arranged on the output side of the internal input/output of the load/store unit 2, more particularly in the data path from the load/store unit 2 to the ALU 5. The reverse path from the ALU 5 to the load/store unit 2 is not affected by the bit pattern generator 6. If there is a data path between the load/store unit 2 and the register unit 4, as indicated by the dashed lines in Figure 8, then the bit pattern generator 6 is also operatively arranged in the data path between those two components. For this optional data path, the operation of the bit pattern generator 6 will be the same as described with reference to the first embodiment and is thus not further discussed. The bit pattern generator 6 is configured to generate as an output a NaN value. In the course of a load completion, a datum is loaded from the load/store unit 2 into a destination operand location of the ALU 5. The bit pattern generator 6 is responsive to the error flag in the load/store unit 2 for the datum being loaded. If the error flag is not set, then the datum is passed unchanged from the load/store unit 2 to the ALU 5, i.e. the bit pattern generator 6 has no operational effect. On the other hand, if the error flag in the load/store unit 2 is set, indicating a failed load, then the bit pattern generator 6 acts to load the reserved NaN value into the destination operand of the ALU, instead of the datum loaded into the load/store unit 2 from the external storage 3.

[0050] The operational effect of the bit pattern generator 6 in the CISC processor of the sixth embodiment is thus generally analogous to that in the RISC proces-

sor of the first embodiment, the difference being that data is loaded into, and downloaded from, the ALU 5 rather than the register unit 4. Both these embodiments have in common that the bit pattern generator 6 is arranged to act as an output stage of the load/store unit located on the processor-internal side of the load/store unit 2.

[0051] Figure 9 shows a processor according to a seventh embodiment of the invention interfaced to an external storage element. The processor unit of the second embodiment includes an instruction unit 1, a load/store unit 2, a register unit 4 including a set of integer registers, and an ALU 5 for integer arithmetic operations. The processor unit is connected to an external storage element in the form of a memory cache 3, the processor unit being connected to the cache 3 through the load/store unit 2. The components of the processor unit are connected generally in the manner of a conventional CISC processor. Namely, there is a line of control from the instruction unit 1 to the load/store unit 2, from the instruction unit 1 to the ALU 5, and also bi-directional lines of control between the load/store unit 2 and the ALU 5 and between the register unit 4 and the ALU 5. Optionally, the bi-directional line of control between the register unit 4 and the ALU 5 may have an extension to the load/store unit 2 in the manner of a RISC processor, as shown by dashed lines in Figure 9.

[0052] By contrast to a conventional CISC processor unit, the communication paths between the load/store unit 2 and the ALU 5 and between the register unit 4 and the ALU 5 are both provided with one bit of extra bandwidth through additional bus lines 7 and 7' respectively. The additional bus lines 7 and 7' are shown in Figure 9, for the sake of convenience only, as being separate from the remaining internal bus lines. An extension of the bus line 7' may optionally be provided to link the register unit 4 to the load/store unit 2, as shown in Figure 9 with dashed lines. The bus line extension would provide a similar functionality to the corresponding connection in the RISC processor of the second embodiment.

[0053] The processor of Figure 9 differs from a conventional CISC design in that the load/store unit 2 and ALU 5 are configured so that a set error flag in the load/store unit will load into a destination operand of the ALU 5 as one or more poison or valid bits. Similarly, poison/valid bits pass with their associated data during communication from the ALU 5 to the load/store unit. Moreover, the ALU 5 is configured to propagate the poison/valid bits present in any operand of an operation to the result or results of that operation, and to return the results together with the propagated poison/valid bits to either the load/store unit 2 or the integer registers in the register unit 4, using the additional bus lines 7 and 7' respectively for communication of the poison/valid bits. In this way, the same functionality as described above for the first and sixth embodiments in respect of floating-point registers, i.e. the quiet propagation of a symbolic

entity for tagging data validity, can also be provided for integer registers.

[0054] Each of the integer registers of the register unit 4 have a poison/valid bit responsive to loads from the ALU 5, and optionally the load/store unit 2, of failed, speculatively loaded, data. The poison/valid bits thus indicate the integrity of the data present as results of an operation in the ALU 5 and data held in the integer registers.

[0055] According to an eighth embodiment of the invention (not separately shown) there is provided a combined integer and floating point CISC processor which is a hybrid of the CISC processors of the sixth and seventh embodiments of the invention in which the register unit 4 has a set of floating-point registers and a set of integer registers, both of which support propagation of a data-integrity signifying symbolic entity. A bit pattern generator 6, as described and illustrated in connection with the sixth embodiment, is provided to support the functionality for the floating-point registers and a bus line 7 and modified internal hardware design of the ALU 5 is provided to support the same functionality for the integer registers, as described and illustrated in relation to the seventh embodiment.

[0056] According to a ninth embodiment of the invention (not separately shown) there is provided a CISC processor with an ALU that supports both floating point and integer arithmetic operations. Floating-point operations are handled in the manner described for the sixth embodiment and integer operations are carried out in the manner described for the seventh embodiment. In addition to these operations, there will also be integer-to-floating-point conversion operations and floating-point-to-integer conversion operations, substantially as illustrated in and described with reference to Figure 7.

[0057] In any of the above-described embodiments of the invention, the ALU 5 may be further modified as follows. Selected types of operation supported by the ALU 5 may, given a certain operand or combination of operands, necessarily cause an invalid or undefined result. One example would be a divide-by-zero operation in integer long-division. Another example would be an overflow in an integer or floating-point arithmetic operation. In such cases, the ALU 5 may be configured to define the result as the reserved QNaN value in the case of a floating-point result or by setting the poison bit in the case of a floating-point or integer result. In either case, the symbolic entity defined by the QNaN value or the poison bit would then be stored in the associated register of the register unit and, if further operations were carried out on that datum, these would be handled in the same way as described above, thus allowing failure testing to be delayed also in the case of ALU generated errors as well as errors resulting from invalid speculative load results.

[0058] Indeed this functionality could be provided on its own without the capability of delaying faults on loads. For example, a further embodiment of the inven-

tion provides a processor for integer operations configured to support poison/valid bits that are propagatable through ALU operations, as described with reference to the second embodiment, wherein the ALU is configured to return a result with a set poison bit (or unset valid bit), when, for one or more prespecified operations, an operand or combination of operands have a predefined value or combination of values. Such a processor may combine floating-point and integer support.

Claims

1. A processor comprising a register unit having a set of registers, an arithmetic logic unit and a load/store unit, the load/store unit being connected to at least one of the register unit and the arithmetic logic unit by a data path and having an error flag for marking as invalid a datum received from an external storage element following a load which has resulted in an error, the processor further comprising a symbolic entity transmitter operatively arranged in said data path from the load/store unit so as to transmit a symbolic entity to said at least one of the register unit and the arithmetic logic unit when an invalid datum is output from the load/store unit for said at least one of the register unit and the arithmetic logic unit, the arithmetic logic unit being configured to propagate the symbolic entity, when present in an operand of an operation carried out by the arithmetic logic unit, to a result of the operation.
2. A processor according to claim 1, wherein the load/store unit is arranged to communicate with the register unit, the symbolic entity transmitter being operatively arranged in the data path from the load/store unit to the register unit.
3. A processor according to claim 1, wherein the load/store unit is arranged to communicate with the arithmetic logic unit, the symbolic entity transmitter being operatively arranged in the data path from the load/store unit to the arithmetic logic unit.
4. A processor according to claim 1, wherein the load/store unit is arranged to communicate with the arithmetic logic unit, the symbolic entity transmitter being operatively arranged in the data path from the load/store unit to the arithmetic logic unit and the register unit.
5. A processor according to any one of claims 1 to 4, wherein the set of registers is a set of floating-point registers and the symbolic entity transmitter includes a bit pattern generator operatively arranged to output a Not-a-Number value when an invalid datum is received from the load/store unit.
6. A processor according to claim 5, wherein the arithmetic logic unit is configured to propagate the Not-a-Number value as a Quiet-Not-a-Number value through operations carried out in the arithmetic logic unit.
7. A processor according to claim 6 responsive to a software command code testing a datum for the presence of a Quiet-Not-a-Number value.
8. A processor according to claim 7, responsive to the software command code so as to cause execution of a branch conditional on the result of the testing, whereby an invalid datum can be handled.
9. A processor according to claim 7, responsive to the software command code conditional on the result of the testing so as to cause conversion of the Quiet-Not-a-Number value to a Signaling-Not-a-Number value, whereby subsequent use of the invalid datum will generate a trap.
10. A processor according to any one of claims 1 to 4, wherein the set of registers includes integer registers, and the symbolic entity transmitter is operatively arranged to output one of poison and valid bits as a symbolic entity, responsive to the error flag in the load/store unit.
11. A processor according to claim 10 responsive to a software command code testing a datum for the presence of one of poison and valid bits.
12. A processor according to claim 11, responsive to the software command code so as to cause execution of a branch conditional on the result of the testing, whereby an invalid datum can be handled.
13. A processor according to any one of claims 1 to 4, wherein the set of registers includes floating-point registers and integer registers, the symbolic entities being Not-a-Number values for floating-point data and one of poison and valid bits for integer data.
14. A processor according to claim 13, wherein the arithmetic logic unit is configured to propagate the Not-a-Number value as a Quiet-Not-a-Number value through operations carried out in the arithmetic logic unit.
15. A processor according to any one of the preceding claims, wherein the set of registers includes floating-point registers and integer registers, the symbolic entities being one of poison and valid bits both for floating-point data and for integer data.
16. A processor according to any one of the preceding claims, wherein the arithmetic logic unit and the register unit are interconnected so as to transmit

and receive the symbolic entities with respective data passing between the arithmetic logic unit, the register unit and the load/store unit.

17. A processor according to any one of the preceding claims; wherein, for at least one prespecified type of operation, the arithmetic logic unit is configured to generate the symbolic entity as a result when an operand of the operation is not an invalid datum and has a predefined value. 5 10
18. A processor according to any one of the preceding claims, wherein, for at least one prespecified type of operation, the arithmetic logic unit is configured to generate the symbolic entity as a result when at least two operands of the operation have a predefined combination of values and are not invalid data. 15
19. A method of operation of a processor comprising an instruction unit, a load/store unit, a register unit comprising a set of registers, and an arithmetic logic unit, the method comprising the steps of: 20
- a) the instruction unit issuing a load request for a datum to an external storage element; 25
 - b) the load being carried out, but returning an invalid datum to the load/store unit;
 - c) the load/store unit setting an error flag for the invalid datum; 30
 - d) the load completing by loading a symbolic entity as at least a part of the datum into one of the register unit and the arithmetic logic unit;
 - e) the arithmetic logic unit carrying out an operation having the datum as an operand such that the symbolic entity associated with the invalid datum is conveyed to a result of the operation; and 35
 - f) outputting from the arithmetic logic unit the result with symbolic entity into one of the register unit and the load/store unit. 40
- 45
- 50
- 55

Fig. 1




Fig. 2



Fig. 3




Fig. 4




Fig. 5




Fig. 6




Fig. 7




Fig. 8






Fig. 9





Europäisches Patentamt
Europ an Patent Office
Office européen des brevets



(11) **EP 1 026 582 A3**

(12) **EUROPEAN PATENT APPLICATION**

(88) Date of publication A3:
04.10.2001 Bulletin 2001/40

(51) Int Cl.7: **G06F 9/38**

(43) Date of publication A2:
09.08.2000 Bulletin 2000/32

(21) Application number: **00300821.6**

(22) Date of filing: **02.02.2000**

(84) Designated Contracting States:
**AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE**
Designated Extension States:
AL LT LV MK RO SI

(72) Inventor: **Harris, Jeremy G.**
Buckinghamshire HP8 4AF (GB)

(74) Representative: **Haines, Miles John et al**
D. Young & Co.
21 New Fetter Lane
London EC4A 1DA (GB)

(30) Priority: **04.02.1999 US 243982**

(71) Applicant: **SUN MICROSYSTEMS, INC.**
Palo Alto, California 94303 (US)

(54) **Handling of load errors in computer processors**

(57) In a RISC or CISC processor supporting the IEEE 754 Not-a-Number (NaN) standard and of the kind comprising a load/store unit, a register unit and an arithmetic logic unit, and wherein the load/store unit has an error flag for marking a datum loaded to the load/store unit following a load which has completed, but resulted in an error, the processor is provided with a bit pattern generator operatively arranged in an output path from the load/store unit to at least one of the register unit and the arithmetic logic unit so that a Not-a-Number value for the invalid datum is loaded into a destination one of the floating-point registers or the arithmetic logic unit. The arithmetic logic unit is configured to propagate the

Not-a-Number value as a Quiet-Not-a-Number (QNaN) value through its operations. The QNaN value may be tested for in a datum by a system software command code provided for that purpose. In other embodiments of the invention, similar functionality is provided for integer registers using poison/valid bits in conjunction with an arithmetic logic unit designed to propagate the poison/valid bits through its operations. An advantage to be gained by this design is that it becomes possible to delay testing the results of a non-faulting load, since the QNaN-like symbolic entity will propagate with the results of operations on an invalid datum, thereby keeping track of the integrity of the data.

EP 1 026 582 A3



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number

EP 00 30 0821

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
A	US 5 692 169 A (WORLEY JR WILLIAM S ET AL) 25 November 1997 (1997-11-25) * abstract * * column 9, line 32 - column 10, line 46 *	1,7-13, 16-19	G06F9/38
A	EP 0 789 298 A (HEWLETT PACKARD CO) 13 August 1997 (1997-08-13) * abstract * * page 7, line 12 - line 40 * * page 8, line 19 - line 40 *	1,7-13, 16-19	
P.A	JP 2000 029866 A (ADVANCED RISC MACH LTD) 28 January 2000 (2000-01-28) & US 6 189 094 A (SEAL DAVID JAMES ET AL) 13 February 2001 (2001-02-13) * column 17, line 40 - line 49 * * column 30, line 5 - line 20 *	1,5,6, 13,14, 17-19	
A	GARCIA L C ET AL: "Storage Access-Exception Detection for Pipelined Execution Units" IBM TECHNICAL DISCLOSURE BULLETIN, IBM CORP. NEW YORK, US, vol. 25, no. 12, May 1983 (1983-05), pages 6711-6712, XP002165704 ISSN: 0018-8689 * the whole document *	1,19	TECHNICAL FIELDS SEARCHED (Int.Cl.7) G06F
A	US 5 799 179 A (EBCIOGLU KEMAL ET AL) 25 August 1998 (1998-08-25) * abstract * * column 3, line 30 - column 4, line 40 *	1,19	
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 10 August 2001	Examiner Moraiti, M
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application I : document cited for other reasons & : member of the same patent family, corresponding document			

FP0 FORM 1503 03 82 (P04C31)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 00 30 0821

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

10-08-2001

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5692169 A	25-11-1997	DE 19534752 A	25-04-1996
		GB 2294341 A	24-04-1996
		JP 8123685 A	17-05-1996
		US 5778219 A	07-07-1998
		JP 5241864 A	21-09-1993
EP 0789298 A	13-08-1997	US 5778219 A	07-07-1998
		JP 9223017 A	26-08-1997
JP 2000029866 A	28-01-2000	US 6216222 B	10-04-2001
		US 6189094 B	13-02-2001
		GB 2339312 A	19-01-2000
		GB 2339623 A	02-02-2000
		JP 11353179 A	24-12-1999
		CN 1303502 T	11-07-2001
US 5799179 A	25-08-1998	CN 1136182 A	20-11-1996
		CZ 9702084 A	17-12-1997
		DE 69600995 D	24-12-1998
		DE 69600995 T	08-07-1999
		WO 9623254 A	01-08-1996
		EP 0804759 A	05-11-1997
		JP 3093624 B	03-10-2000
		JP 8263287 A	11-10-1996
		PL 321542 A	08-12-1997

EPO FORM 2459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

This Page Blank (uspto)